The Interplay of Computer Power, Computer Architecture and Numerical Algorithms in the progress of Numerical Weather Prediction

Anthony Hollingsworth
European Centre for Medium-Range Weather Forecasts
Shinfield Park, Reading, RG2 9AX, UK

Given the importance of numerical accuracy in the data assimilation and modelling aspects of Numerical Weather Prediction (NWP), it is a perennial truth that the current need for computer power far exceeds the current ability of the computer industry to satisfy that need, at the current level of funding.

The NWP community has used several strategies to reduce this perennial gap between computer needs and computer funding.

The first strategy has been to demonstrate to the public, to industry and to politicians the value (in terms of life and property) of sustained investment in NWP. The success of NWP in delivering increasingly useful forecasts, along with the increasing vulnerability of the population to weather-related hazards, have together ensured the maintenance in most countries of the annual spend on observations (both satellite and in-situ), on technical infrastructure (high-performance computing, data-handling capability, tele-communication) and on the human resources needed to deliver improved forecasts. In parallel, over many years, the computer industry has been able to double its capability every 18-24 months at more or less constant cost. Thus a sustained annual spend on computing by NWP institutes can make the benefits of Moore's Law available to the NWP community to provide sustained improvements in forecast accuracy.

A second NWP strategy to reduce the gap between computer needs and available resources has been to ensure that operational codes are adapted to run efficiently on the widest possible range of current and emerging computer architectures. If an NWP institute's operational codes are designed to be flexible, parallel and portable enough to run efficiently on existing and emerging architectures and if the operational NWP production-system is designed to facilitate a straight-forward migration from one architecture to another, then two benefits are available. The first is a mutually beneficial dialogue between NWP institutes and vendors on issues of hardware and software design and upgrades. The second benefit, for NWP, is to maximise competition for NWP's business, which delivers real benefits in the amount of computing realised per euro or per dollar spent.

A notable example of this strategy during the 1990s was the work of the RAPS consortium (Hoffmann pers. comm.. 1992) where RAPS stands for Real Applications of Parallel Systems. In the RAPS initiative a group of leading computing labs in Europe got together with leading computer manufacturers to address the problems of exploiting new computer architectures in a number of very different applications, including NWP and climate simulation. The fact that the RAPS consortium has flourished for 15 years (since 1989) and has been jointly funded throughout by the computer manufacturers and by the applications centres, is clear confirmation that there is a strong mutual benefit from the pre-benchmarking activity.

The third strategy used by NWP centres to reduce the gap between resources and needs is to improve the efficiency of the numerical algorithms used in the data assimilation system and in the model integrations.

For example, the development of the incremental formulation for the four-dimensional variational assimilation system (4D-Var) enabled a much-earlier-than expected operational implementation of this most effective method to use observations (which are the largest single cost element of NWP). Other examples in the assimilation area include the recent development of effective pre-conditioning methods to accelerate the convergence of 3D-Var and 4D-Var.

A prime example of an efficiency gain in the modelling area was the development (in the 1960s) and widespread adoption (in the 1970s) of semi-implicit integration schemes. These schemes provided considerable gains in the efficiency of NWP model integrations, equivalent to at least a new generation of super-computers.

Even more dramatic gains have been achieved in the 1990s with the development of semi-Lagrangian time-integration schemes, which impact the costs of both model and assimilation system. For example, with its present computer resources and operational time slot, and with the time-integration schemes used up to 1990, ECMWF could produce at best a 3D_Var analysis and a 2 day forecast at the current operational resolution of T511. Thanks to the efficiency gains offered by the semi-Lagrangian scheme, the same computer resources can produce daily a more accurate (and expensive) 4D-Var analysis, a 10-day forecast at T511 (40 km) , and a 51 member ensemble at T255 (80km). The efficiency gains more than pay for the improved forecasts model, for the more accurate assimilation system and for the ensemble forecast system; the latter plays an increasingly important role in identification and management of weather risk in many industries and applications.

The presentation will review the interplay of numerical and computing developments in recent years, and consider some of the issues to be addressed in the development of operational Earth-system model/ assimilation capabilities.